GARA: An Architecture for Advanced Reservations

Alain Roy, Linda Winkler
roy@mcs.anl.gov, winkler@mcs.anl.gov
Argonne National Laboratory
7/21/99

Goals

- Provide end-to-end Quality of Service to applications. This requires:
  - Discovery and selection of resources
  - Allocation of resources
  - Advance reservation of resources

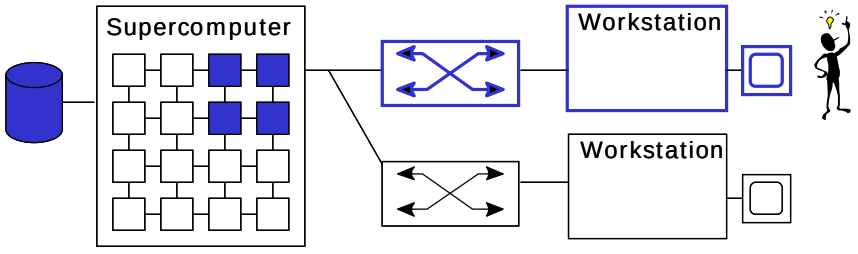

Difficulties/Solutions

- Lack of support for advanced reservations
  - We can use existing advanced reservation mechanisms if available or supply our own
- Heterogeneous resources
  - We provide uniform interfaces
- Need to work with complex sets of resources
  - We use co-reservation and co-allocation agents
- Resources in different administrative domains
  - We use the Globus security infrastructure

Solution: GARA

Globus Architecture for Reservation and Allocation

- Three important contributions:
  - Advance reservations are first-class entities
  - Uniform treatment of underlying resources (CPU,network, disk, etc.)
  - Layered architecture enables generic co-reservation and co-allocation agents

GARA Basics—Reservations

- There is a generic “reservation”, which has several properties:
  - Start Time (“now” or future) and Duration
  - Resource type/Underlying resource identifier
  - Resource-specific (bandwidth, % CPU…)
- All reservations are treated uniformly:
  - Create/Modify (Given properties)
    => Returns Reservation Handle
  - Destroy
  - Monitor (Callbacks or Polling)

A picture of what happens (creating a reservation)

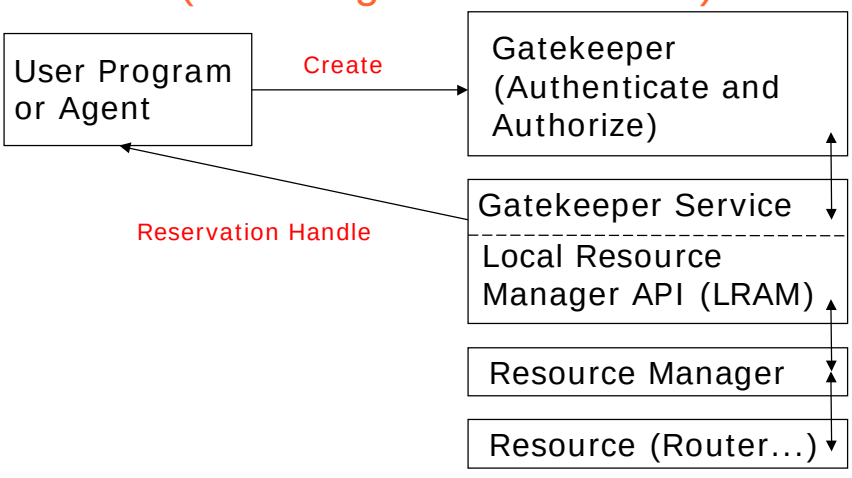

The Resource Manager

- To track advanced reservations, we use a "slot table"
- We assume exclusive access to resource through the resource manager

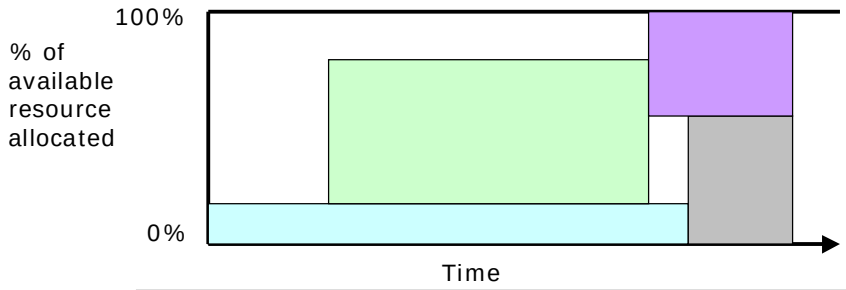

The Big Picture

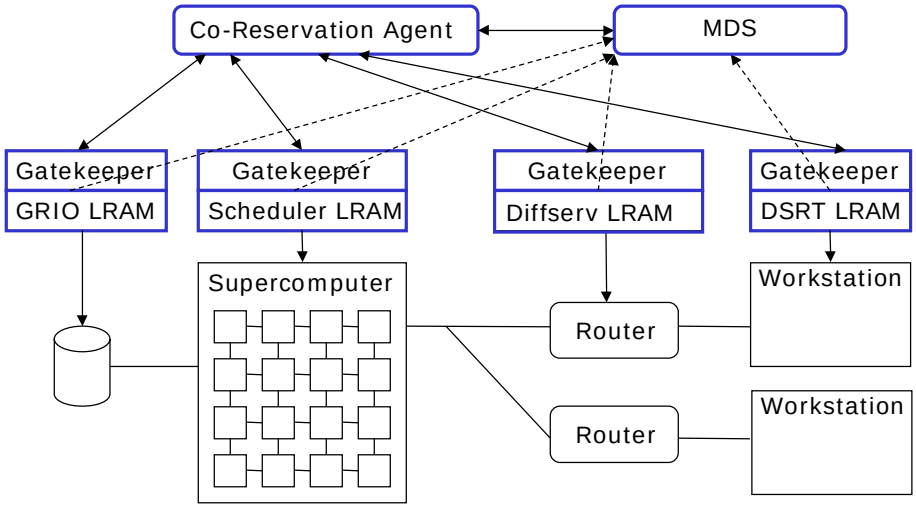

End-to-End Network Reservations

- Algorithm for a single network reservation
  - Discover which gatekeepers/resource managers need to be contacted (MDS)
  - Make reservation at each one.

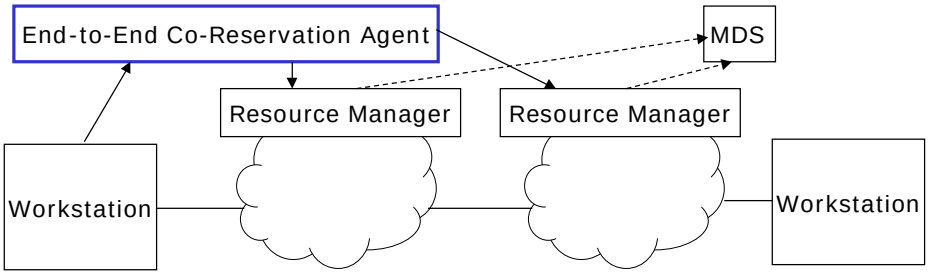

Current Status

- A working prototype of GARA exists
  - Differentiated Services-like QoS
  - Real-Time CPU Scheduling
  - DPSS Disk Access
- An early prototype of the end-to-end reservation agent exists
- We have done some experiments on a testbed
- Working with bleeding-edge adopters has started
- Adding WWW-based calendar interface

WWW-based Calendar Interface

Future Work

- Supporting more resource types
  - Job Schedulers
  - Alternative Network QoS (RSVP)
  - Disk (GRIO)
  - Others (Memory, etc…)
- Experimentation with real applications
- High-level agents to simplify usage
- Deployment with Globus